# When ADC commit changes to the AD

By Marcelo Calbucci

Keywords: ADC, USNChanged, msExchServer1HighestUSN, msExchServer2HighestUSN, highestUSNCommited

The Active Directory Connector is a long-run process. It is a service that runs under Windows 2000, and it will only stop when you shutdown the machine (or if you choose to manually stop it). For this reason, it caches a lot of information that it uses on memory. One of this information is the Highest USN.

ADC uses the Highest USN to determine the next search it will do. So, if your Connection Agreement starts running when the 'highestCommittedUSN' was 10,000, we will know that at the end of the first cycle, we read all entries that had their uSNChanged less than 10,000. On the next cycle, we'll read the 'highestCommittedUSN' (let's say 15,000), and read all entries with uSNChanged greater than 10,000, and at the end we know that the highest USN visited was 15,000, and so on.

The problem of keeping this information only in memory is that when the machines shutdown, or the service is stopped, we will lose it.

For this reason, we use some attributes on the Connection Agreement to store this data. These attributes are: msExchServerXHighestUSN, msExchServerXHighestUSNVector and msExchServerXLastUpdateTime, where X is 1 or 2.

The msExchServerXLastUpdateTime is not used by the ADC; it is just stamped for reference on the Connection Agreement. This attribute may not contain accurate information of when was the last time the CA run. It is there for backward compatibility reasons.

The msExchServerXHighestUSNVector (notice the Vector at the end) will be present only if the Server X is a Windows 2000 server. It will be used if a CA was replicating from one DC, but them you point it to another DC. This prevents ADC from missing entries, because we keep a list of "highestCommittedUSN" on all DCs on your organization.

Finally, the msExchServerXHighestUSN is the attribute that ADC will decide from where to start searching the Directory. If this attribute is set to zero, it has the same effect of setting "msExchDoFullReplication" to TRUE, or to right click a Connection Agreement and choose on the "Schedule" page of properties to "Replicate the entire directory the next time the agreement run." All this solutions will make ADC start from the scratch and replicate

everything again, except that for the msExchServerXHighestUSN you can specify only one direction of the replication.

Use msExchServer1HighestUSN to specify the Active Directory and msExchServer2HighestUSN to specify the Exchange Directory.

We update, in memory, the Highest USN every time we finish a cycle or a search block (see the KB “ADC Search Block”), and we’ll commit the changes to the Active Directory every 24 hours if this is not a new Connection Agreement. If it is a new CA, then we’ll update the AD every 30 minutes. This is because the first time you set a CA, it has to do a lot of additions, which are expensive operations, so we want to be sure that even if you have a power outage in the middle of a cycle, you won’t loose more than 30 minutes of replication.

We will also immediately commit to the AD in more three special situations. First, if you’re shutting down the ADC Service, we need to write the information that we have in memory to the AD, so ADC doesn’t loose it. Second, when it is the end of the first cycle of a new Connection Agreement. This is to mark that we already did all the first-time adds operations. And third, when you’re re-homing your CA to run on a different server. We need this because RAM is not shared between servers, so you need to commit it to the AD that it is shared resource between servers.

Let’s illustrate this process with an example.

Imagine that you have your Exchange server A and your Windows server B. This is the first time that you’re running a Connection Agreement between them. Most likely you’ll have an Exchange server full of Mailboxes, Custom Recipients and Distribution Lists. Let’s also imagine that you have a total of 20,000 entries on this server.

Basically, what ADC needs to do is to sync all those 20,000 entries from server A to server B, and sync the entries from B to A. This is probably going to take a while to be done, and it will depend mostly on the performance of both servers and the performance of the network. For this example, let’s also say that the highestCommittedUSN on server A is 80,000, since you probably do more modifications then additions.

The ADC will start replicating this Connection Agreement and it will read the ‘highestCommittedUSN’, which is 80,000, and the last USN when we ran the CA, which is zero. So the first search ADC will do is for the uSNChanged between 0 and 10,000 on the Server A. After reading and updating all the entries on this block, it will store in memory that it visited the uSNChanged up to the 10,000. If it is passed 30 minutes since the when we started, we “record” this information to AD, otherwise we keep it in memory only. Then ADC will read the next block from 10,001 to 20,000. After that it will commit the change to AD if it passed 30 minutes from the last time we

committed, and so on, until we finish the first cycle for this CA on both directions, on which case we write to the AD the msExchServer2HighestUSN equal to 80,000 and the value of the highestUSNCommited on the AD to the attribute msExchServer1HighestUSN.

After this first time, every time the CA runs it will only update the CA entry every 24h, even if you start and stop the service.

So, what you have is that on the worst case scenario, if powers goes down, and you’re doing the first replication ever, you’ll loose no more then 1h of replication.